(U) Dynamic Distributed Cooperative Control of Multiple Heterogeneous Resources

Dr. Michael J. Hirsch
Raytheon Company
Orlando, FL
Michael_J_Hirsch@Raytheon.com

Daniel Schroeder
Raytheon Company
State College, PA

ABSTRACT

This research is concerned with dynamically determining appropriate flight patterns for a set of autonomous UAVs in an urban environment, with multiple mission goals. The UAVs are tasked with searching the urban region for targets of interest, and tracking those targets that have been detected. We assume that there are limited communication capabilities between the UAVs, and that there exist possible line of sight constraints between the UAVs and the targets. Each UAV (i) operates its own dynamic feedback loop, in a receding horizon framework, incorporating local information (from the perspective of UAV i) as well as remote information (from the perspective of the 'neighbor' UAVs) to determine the tasks to perform and the optimal trajectory of UAV i (and neighbor UAVs) over the planning horizon. This results in a decentralized and more realistic model of the real-world situation. As the coupled task assignment and flight route optimization formulation is *NP-hard*, a hybrid heuristic for continuous global optimization is developed to solve for the flight plan and tasking over the planning horizon. Metrics capturing the price of anarchy and price of decentralization are developed, and experimental results are discussed.

Keywords: Cooperative Control, Decentralized Control, Price of Anarchy

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
OCT 2012	**N/A**	**-**

4. TITLE AND SUBTITLE	
Dynamic Distributed Cooperative Control of Multiple Heterogeneous Resources	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Raytheon Company Orlando, FL**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM202976. 2012 Joint Meeting of the Military Sensing Symposia (MSS) held in Washington, DC on October 22-25, 2012.

14. ABSTRACT
This research is concerned with dynamically determining appropriate flight patterns for a set of autonomous UAVs in an urban environment, with multiple mission goals. The UAVs are tasked with searching the urban region for targets of interest, and tracking those targets that have been detected. We assume that there are limited communication capabilities between the UAVs, and that there exist possible line of sight constraints between the UAVs and the targets. Each UAV (i) operates its own dynamic feedback loop, in a receding horizon framework, incorporating local information (from the perspective of UAV i) as well as remote information (from the perspective of the neighbor UAVs) to determine the tasks to perform and the optimal trajectory of UAV i (and neighbor UAVs) over the planning horizon. This results in a decentralized and more realistic model of the real-world situation. As the coupled task assignment and flight route optimization formulation is NP-hard, a hybrid heuristic for continuous global optimization is developed to solve for the flight plan and tasking over the planning horizon. Metrics capturing the price of anarchy and price of decentralization are developed, and experimental results are discussed.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**SAR**	**14**	

1.0 Introduction

Military forces face ever increasing challenges to provide timely and accurate information on targets of interest, especially those targets that are mobile and elusive in nature. These tasks expand significantly in complexity when the targets are operating in an urban environment. The use of multiple unmanned vehicles to provide this target information allows our military personnel to stay out of the line of fire. However, many remotely controlled UAVs (i.e., swarms) require as many skilled pilots as there are swarm members, and these pilots must be able to deconflict airspace demands, mission requirements, and situational changes in near real time [3]. On the other hand, autonomous unmanned vehicles allow military personnel to focus on more important issues like interpreting the gathered information, as opposed to determining how to acquire the information [4]. Hence, there is a need to build intelligent unmanned vehicles that can plan and adapt autonomously to the environment they perceive, while also collaborating with the human-in-the-loop as appropriate [46]. The clear benefit is shortened mission-critical decision chains.

There has recently been much research done in the area of autonomous vehicle control for surveillance type missions. Almost all of the research has dealt with centralized cooperative control, with little research concerned with the more realistic decentralized problem [5]. Steinberg [6] provided an overview on research and limitations of autonomous technologies for the control of heterogeneous unmanned naval vehicles. Experiments in this paper examined aspects such as multi-vehicle task allocation, dynamic replanning of vehicle tasks, as well as human-in-the-loop management. Constraints considered in the experiments included pop-up threats, adverse weather conditions, and communication issues between the autonomous vehicles, among others. Ahmadzadeh et al. [8] described their Time Critical Coverage planner as a component of the Office of Naval Research Autonomy program, ICARUS. Each autonomous vehicle was modeled as a Dubin's vehicle [9], where-by the vehicles were assumed to be point masses with constant speed and a prescribed minimum turning radius. The vehicles also had prescribed starting and ending spatial-temporal locations, as well as polygonal obstacles to be avoided throughout flight. The objective was to determine the flight path of the UAVs to maximize the total sensor footprint over the region of interest. The algorithm utilized to solve this problem was based on sampling a discretized search graph [10]. Shetty, Sudit, and Nagi [13] considered the strategic routing of multiple unmanned combat vehicles to service multiple potential targets in space. They formulated this as a mixed-integer linear program, and through a decomposition scheme looked at solving the target assignment problem (vehicles to targets) and then determining the tour that each vehicle should take to service their assigned targets (a classical vehicle routing problem). They implemented a tabu-search heuristic to find solutions to their problem. However, they assumed the vehicles were holonomic, which enabled the mixed-integer linear program formulation. Schumacher and Shima [16] considered the problem of wide area search munitions, which are capable of searching for, identifying, and attacking targets. Whenever a new target is found, or a new task needs to be assigned, a capacitated transshipment assignment problem is solved, to determine the optimal assignment of munitions to tasks. Note that from one solution to the next solution, the assignment can change significantly. Schouwenaars et al. [21] considered fuel-optimal paths for multiple vehicles. They formulated the problem as a mixed-integer linear program. The vehicles needed to move from an initial to final state, while avoiding vehicular collisions, as well as stationary and moving obstacles. Obstacle positions were assumed known *a priori*. Kingston and Beard [22] presented an algorithm to keep moving UAVs equally spaced (angularly) about a stationary target. The UAVs adapted their spacing based on local communication with other UAVs. Velocity bounds were derived so that the UAVs satisfied their kinematic constraints. They extended this approach to the case when the target was moving, but the path and velocity of the target was assumed known throughout. Gu et al. [23] discussed the problem of cooperative estimation using a team of UAVs. They considered the goal of organizing the teams configuration to achieve optimal position and velocity estimates of a moving ground target.

Close range maneuvering and following of a moving target continue to pose significant research challenges [27]. Close-range maneuvering requires real-time dynamic replanning and decision-making, as well as optimization of many parameters (taking into account the physical constraints of the vehicle). The problem becomes even more complex when the goal is to track a moving target for which the target dynamics is not known. In [31, 32, 33], Hirsch et al. developed a decentralized approach for the autonomous cooperative control of UAVs, with the goal of persistent and accurate tracking of moving ground targets in an urban domain. The UAVs were able to share limited information with neighboring UAVs (other UAVs in their communication region), and had to dynamically re-plan their flight paths, incorporating predicted target movements and re-planned flight paths of their neighbors into their own decision making process. Line of sight also plays a big role in the flight plan of the UAVs and was incorporated into the problem through the use of Plüker coordinates [1, 2]. However, it was assumed that the targets were already known. In [44], Hirsch and Schroeder extended the work of [31, 32] to include, in a decentralized framework, the UAVs actively searching for and detecting the targets, before tracking of the targets can commence. Since the UAVs have no knowledge of how many targets are present in the environment, at each decision making step, the UAVs needed to determine both the task to perform (either searching for new targets or tracking those targets already detected) as well as the trajectory to optimally accomplish their task. This resulted in a coupled task assignment / route planning mathematical formulation, which is highly nonlinear. In this research, we look at the problem in [44] when the UAVs have two different types of passive sensors on-board. In addition, we analyze the results using 'price of anarchy' [45] and 'cost of decentralization' metrics.

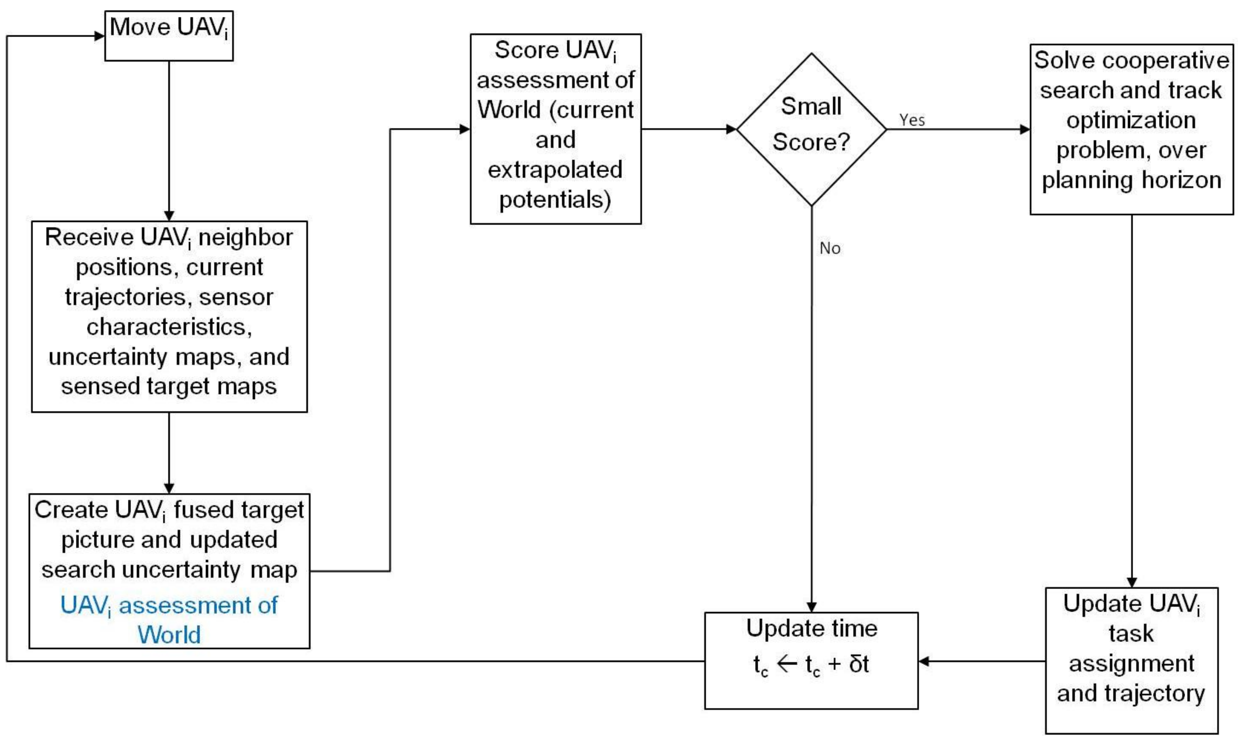

Figure 1: Dynamic feedback loop.

2.0 Problem Description

The overall problem addressed here is to provide (i) continual surveillance over a region of interest and (ii) accurately track all targets detected in this region. Hence, there are two tasks that a UAV can be performing at any given time: searching for targets or tracking those targets already detected. For this problem, there are $\boldsymbol{N}$ heterogeneous autonomous vehicles flying at a fixed altitude. The vehicles are modeled as non-holonomic point masses on a two-dimensional plane with a minimum turning radius (i.e., a Dubin's vehicle [9]). The vehicles have minimum and maximum speed restrictions, as well as a maximum communication range, beyond which they cannot share information. Each vehicle has an independent internal representation of the surveillance region as a continuous search uncertainty map. For each location $\boldsymbol{Y}$ in the region, the search uncertainty map for a particular UAV stores the uncertainty (from the UAVs perspective) of a target currently at location $\boldsymbol{Y}$. While in search mode, the UAV can update its search uncertainty map using information from its own on-board sensor, as well as information from other UAVs currently in search mode that can communicate with this UAV (i.e., its 'neighbor' UAVs). While in tracking mode, the UAV can only update its search uncertainty map using information from neighbor UAVs that are in search mode. Similarly, each UAV has its own internal representation of the targets in the region of interest. When in search mode, the UAV can only update its target map from those neighbor UAVs in track mode. When in track mode, the UAV can update its target map based upon its local sensors, as well as information provided from neighbor UAVs in track mode. The one exception to this is if in search mode, it is possible (and hoped) that a UAV will detect a 'new' target. When this happens, this target is added to the UAVs stored target map, and shared with neighbor UAVs.

The vehicles operate in a decentralized manner; at each time step, every UAV processes an equivalent version of the same dynamic feedback loop, as presented in Figure 1. We discuss the feedback loop from the standpoint of UAV i. At the current time, t_c, the first step is for UAV i to move according to its planned route, and continue to execute its current task (e.g., searching or tracking). Using on-board sensors, UAV i either searches its current field of view for new targets or attempts to sense the targets (known to UAV i) that are within its line of sight. The neighbors of UAV i send information relating to their (a) current positions, trajectories, and tasking; (b) sensor characteristics; (c) search uncertainty maps; and (d) sensed target maps. At this point, UAV i updates its assessment of the world, represented as a fused search uncertainty map and a fused target location map. This fused picture is scored in some qualitative manner, based upon the current positions of UAV i and neighbor UAVs, their current planned trajectories and tasking, as well as the extrapolated search uncertainty map and target tracks, over the planning horizon. If the resultant score is small, or if the elapsed time since the last replan for UAV i is large, then UAV i solves a coupled tasking / route planning optimization problem for itself and its neighbor UAVs, over the planning horizon (discussed in detail in [44]). The solution to this problem is a tasking for each UAV and route for each UAV to fly in order to accomplish their tasking. We note that while UAV i updates its own flight path and tasking, it does not share any of this computed information with its neighbor UAVs. The reason for this is that the neighbors are going through this same feedback loop, and may be incorporating additional knowledge not available to UAV i. The time is then incremented, and the dynamic loop continues.

3.0 Heuristic Approach

To efficiently solve the optimization problem discussed above, we have developed a hybrid heuristic combining Greedy Randomized Adaptive Search Procedures (GRASP) and Simulated Annealing (SA). GRASP is a multi-start local search procedure, where each iteration consists of two phases, a construction phase and a local search phase [38, 39, 40]. In the construction phase, interactions between greediness and randomization generate a diverse set of quality solutions. The local search phase improves upon the solutions found in the construction phase. In our implementation, we have only made use of the construction portion of GRASP. Simulated annealing is a heuristic to find good-quality solutions to optimization problems by approximating the cooling process of metals [41]. At each step, a current solution is perturbed. If the perturbation results in a better solution, then the current solution is replaced. If the perturbed solution is worse than the current solution, it still might replace the current solution; replacement will occur with a probability based on the distance between the current and perturbed solution values and the current temperature in the annealing process. As the heuristic progresses, the temperature is lowered, making it more and more unlikely to replace the current solution with a worse solution.

For our hybrid GRASP-SA heuristic, the GRASP heuristic determines the task assignment for the UAVs, interacting directly with a Simulated Annealing heuristic which determines the routes that the UAVs should fly to accomplish their tasks. The best solution over all of the GRASP multi-start iterations is retained as the final solution. We note here that a solution, for a given UAV, is the tasking and flight path for that UAV, as well as all neighbors of that UAV, over the planning horizon.

4.0 Experiments and Analysis

Figure 2 shows our experimental environment, a simulated urban environment, displayed from an overhead perspective. The yellow rectangles represent buildings. Small white squares represent UAVs; the blue rectangular box emanating from each UAV is the sensor footprint, and the lines extending from each UAV represent its projected flight path. Sensors onboard the UAV can rotate independent of the UAV's heading. The UAVs can communicate with each other, as long as a building doesn't block their line of sight. The segments of the flight paths are colored either pink (used to represent timesteps where the UAV is performing a "search" task) or green (where the UAV is performing a "tracking" task). The small pink squares represent targets. The UAVs have no *a priori* knowledge of the number of targets or their movement dynamics. These targets can be stationary, mobile, or alternate between stationary and mobile at different times.

The colored background in Figure 2 represents the minimum search uncertainty values across all UAVs. The colors range from blue to red: blue representing areas of low search uncertainty, and red representing areas of high search uncertainty. A UAV that performs a search task over a particular area reduces the uncertainty, changing the colors from red to blue. If no UAV has searched over an area for a length of time, the uncertainty increases, and the colors change from blue to red.

Figure 3 depicts time-step 3, when both UAV_1 and UAV_2 detect targets; however, their current trajectories are such that they will overfly these targets, and not be able to track them. UAV_3 is a neighbor of both UAV_1 and UAV_2. Hence, UAV_3 is made aware of these detected targets. As seen in Figure 4, at time-step 4 UAV_3 switches from a search task to a track task, to track the target detected by UAV_1. At time-step 5 (Figure 5), UAV_2 is seen to shown modifying its flight path, so that it can begin tracking the target it detected at time-step 3.

At time-step 35 (Figure 6), UAV_1 detects two targets, and UAV_2 detects a third target. Because these UAVs are neighbors, they communicate all detections to each other. These UAVs switch from searching to tracking the targets at time-step 36 (Figure 7), with modified trajectories to allow them to track all three targets. Both UAVs successfully track these targets until time-step 44, when they switch back to search tasks.

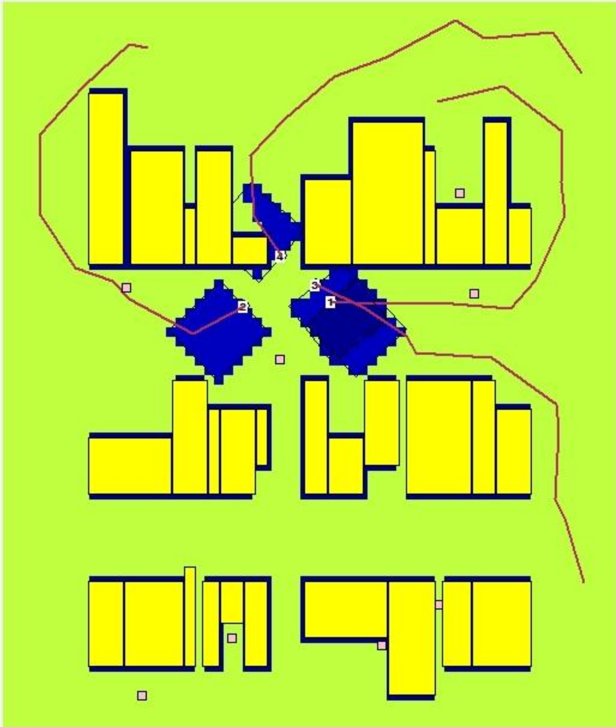

Figure 2: Initial time-step of experiment.

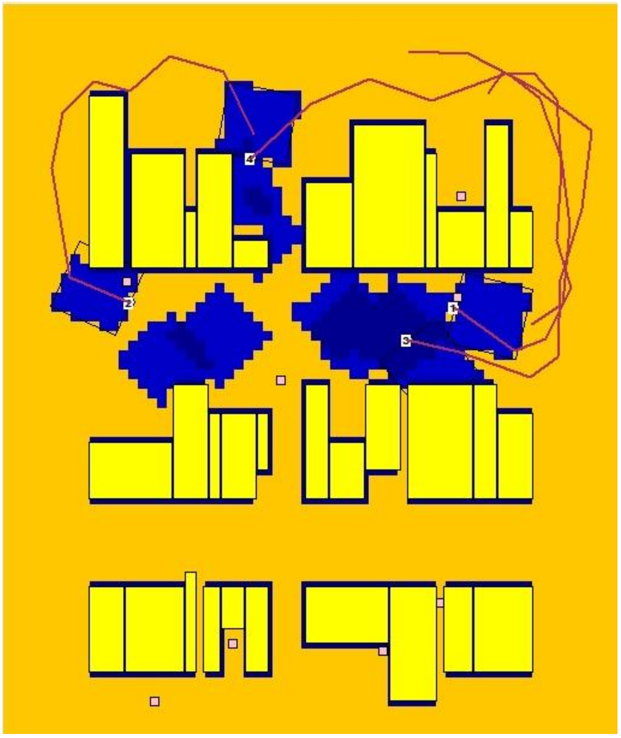

Figure 3: Time-step 3. Both UAVs 1 and 2 detect targets.

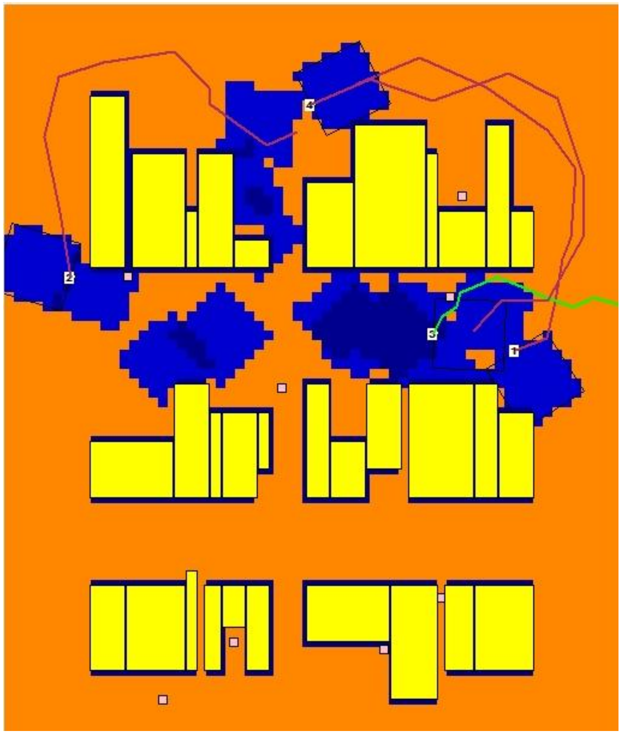

Figure 4: Time-step 4. UAV_3 begins tracking target detected by UAV_1.

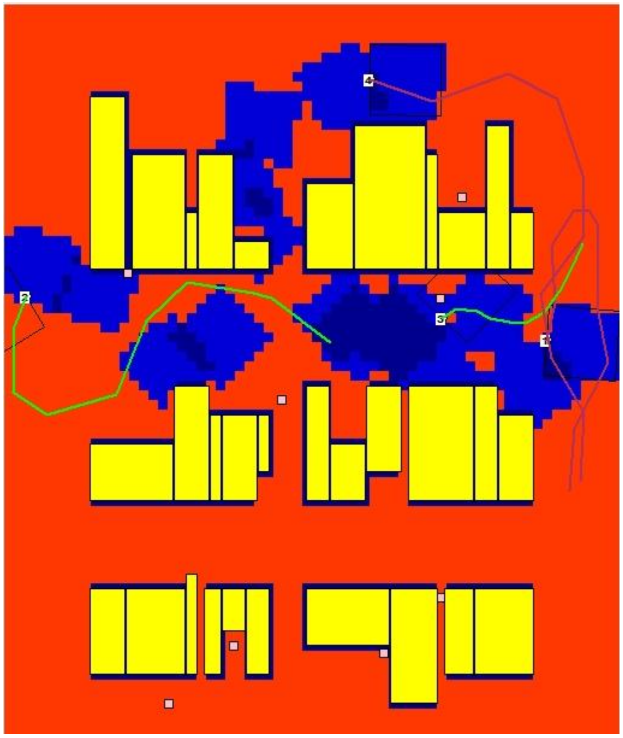

Figure 5: Time-step 5. UAV_2 alters trajectory try to reacquire the target it detected at time-step 3.

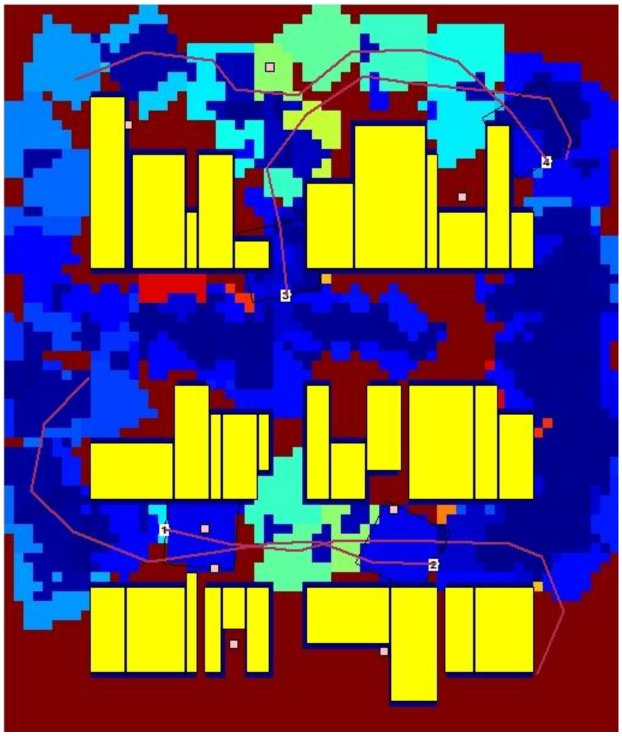

Figure 6: Time-step 35. UAV_1 detects two targets, and UAV_2 detects a third.

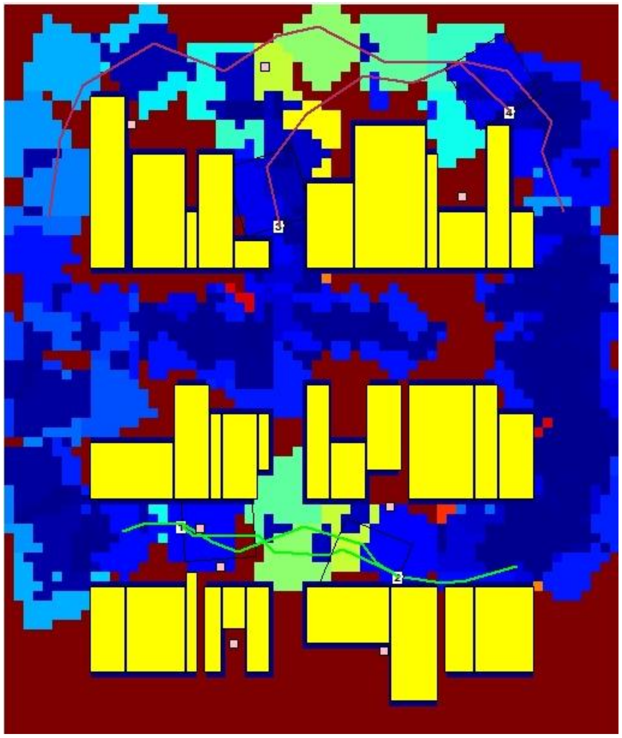

Figure 7: Time-step 36. UAVs 1 and 2 switch from a search to a track task, to keep track of the three targets they detected at time-step 35.

We analyze our experiment using two metrics. The Cost of Decentralization is defined as a measure of how the solution quality degrades from a centralized to a decentralized environment. At each time-step in our experiment, we determined the situation awareness picture of each UAV, and appropriately combined these pictures to create the situation awareness picture for a centralized problem. We then used that centralized situation awareness picture to solve the coupled tasking / trajectory optimization problem for all the UAVs. We computed the number of UAVs where the tasking is different from the centralized to the decentralized solution. That result is shown in Figure 8. For those UAVs with the same tasking, we computed the average distance between their computed trajectories. This is shown in Figure 9. One would rightly expect that a centralized solution would be closer to optimal (in a global sense) than a solution determined in a decentralized environment.

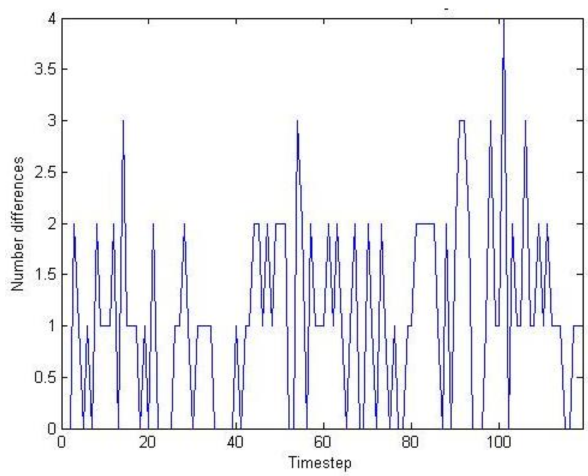

Figure 8: Cost of Decentralization. Number of UAVs with different tasking.

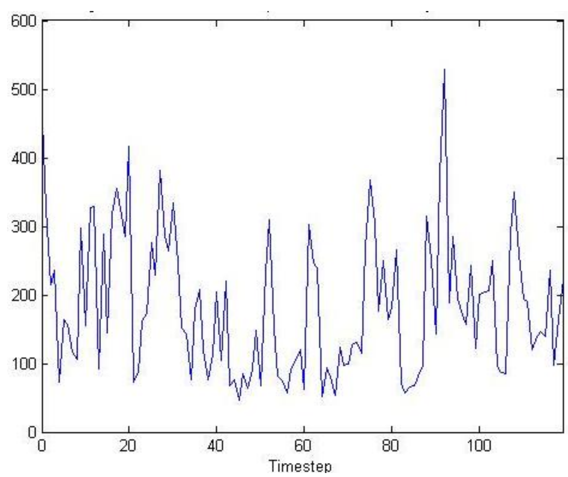

Figure 9: Cost of Decentralization. Average distance (in meters) between trajectories of those UAVs with the same tasking.

The average distance is smallest at time-step 84. Comparing the centralized with the decentralized solution, two of the UAVs have the same tasking and two have different tasking. The two with the same tasking (tracking) have the same tracking situational awareness picture in both the centralized and decentralized environments, and result in very similar computed trajectories. The average distance is largest at time-step 93. At this time-step, only one UAV has the same tasking when comparing the decentralized and centralized solutions. Based on the situational awareness picture in the centralized case, the UAV is routed in the opposite direction as from the decentralized case, resulting in the large average distance.

The price of anarchy metric is adapted from routing in networks [45]. This is defined as the comparison in solution quality between a selfish versus a cooperative approach. At each time-step of the experiment, going through the dynamic feedback loop of Figure 1, each UAV has the possibility of solving a coupled task assignment and routing optimization problem. In the cooperative case, as discussed above, the UAVs solve this problem for themselves and their neighbors. In the selfish case, while each UAV makes use of the information provided by their neighbor UAVs to create their fused situational awareness picture, each UAV solves the optimization problem only for themselves. One would expect that a cooperative solution would be more optimal (in the global sense) than a selfish solution. Again, we computed the number of UAVs where the tasking is different from the cooperative to the selfish solution. That result is shown in Figure 10. For those UAVs with the same tasking, we computed the average distance between their computed trajectories. This is shown in Figure 11.

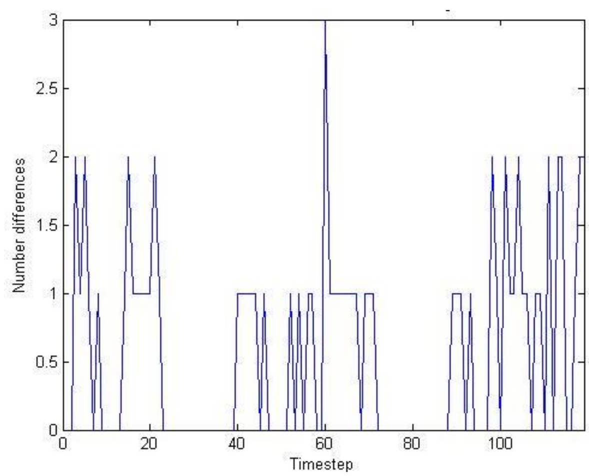

Figure 10: Price of Anarchy. Number of UAVs with different tasking.

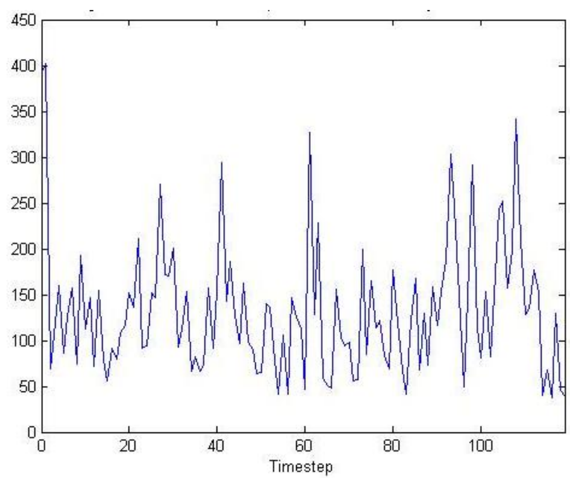

Figure 11: Price of Anarchy. Average distance (in meters) between trajectories of those UAVs with the same tasking.

The smallest average distance occurred at time-step 55. At this time-step, none of the UAVs are in communication with each other, so the cooperative approach and the selfish approach should be the same. Due to stochasticity in the heuristic, one of the UAVs was tasked differently between the cooperative and selfish solutions, and for the three UAVs tasked the same, the trajectories are almost identical.

5.0 Conclusions and Future Research Directions

In this research, we have considered UAVs tasked with searching an urban environment for targets of interest, and tracking those targets that have been detected. The UAVs operate cooperatively, in a decentralized framework. We have also introduced two metrics. The first, the Cost of Decentralization, measures the solution degradation in a decentralized framework versus a centralized framework. The second, the Price of Anarchy, measures the solution degradation as the UAVs move from a cooperative approach to a selfish approach. Future research will include investigating how to modify the tasking for the UAVs on a sub-interval of the planning horizon, to account for collection opportunities, without degrading the original tasking for the UAVs.

6.0 Acknowledgements

We gratefully acknowledge the financial support and technical guidance of Dr. John Zolper and Mr. David Ii, both of Raytheon Company.

7.0 References

1. Semple, J. G. and Kneebone, G. T., Algebraic Projective Geometry, Oxford University Press, 1952.

2. Stolfi, J., Oriented Projective Geometry: A Framework for Geometric Computations, Elsevier, 1991.

3. Bamburger, R. J., Watson, D. P., Scheidt, D. H., and Moore, K. L., "Flight demonstrations of unmanned aerial vehicle swarming concepts," Johns Hopkins APL Technical Digest, Vol. 27, No. 1, 2006, pp. 41–55.

4. Hirsch, M. J., Ortiz-Pena, H., Sapankevych, N., and Neese, R., "Efficient Flight Formation for Tracking of a Ground Target," Proceedings of the National Fire Control Symposium, 2007, pp. 1–16.

5. Shima, T. and Rasmussen, S., UAV Cooperative Decision and Control, SIAM, 2009.

6. Steinberg, M., "Intelligent autonomy for unmanned naval vehicles," Proceedings of SPIE: Unmanned Systems Technology VIII, edited by G. R. Gerhart, C. M. Shoemaker, and D. W. Gage, Vol. 6230, 2006, pp. 1–12.

7. Liang, Z., Chaovalitwongse,W. A., Rodriguez, A. D., Jeffcoat, D. E., Grundel, D. A., and O'Neal, J. K., "Optimization of Spatiotemporal Clustering for Target Tracking from Multi-Sensor Data," IEEE Transactions on Systems, Man, and Cybernetics (C), Vol. 40, No. 2, 2010, pp. 176–188.

8. Ahmadzadeh, A., Buchman, G., Cheng, P., Jadbabaie, A., Keller, J., Kumar, V., and Pappas, G., "Cooperative Control of UAVs for search and coverage," Proceedings of the AUVSI Conference on Unmanned Systems, 2006.

9. Dubins, L. E., "On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents," American Journal of Mathematics, Vol. 79, No. 3, 1957, pp. 497–516.

10. LaValle, S. M., Planning Algorithms, Cambridge University Press, 2006.

11. Ma, C. S. and Miller, R. H., "Mixed integer linear programming trajectory generation for autonomous nap-of-the earth flight in a threat environment," Proceedings of the IEEE Aerospace Conference, 2005, pp. 1–9.

12. ILOG CPLEX, "http://www.ilog.com/products/cplex," Accessed October 2008.

13. Shetty, V. K., Sudit,M., and Nagi, R., "Priority-based assignment and routing of a fleet of unmanned combat aerial vehicles," Computers and Operations Research, Vol. 35, No. 6, 2008, pp. 1813–1828.

14. Kenefic, R. J., "Finding good Dubins tours for UAVs using particle swarm optimization," Journal of Aerospace Computing, Information, and Communication, Vol. 5, 2008, pp. 47–56.

15. Kennedy, J. and Eberhart, R., "Particle Swarm Optimization," Proceedings of the IEEE International Conference on Neural Networks, 1995, pp. 1942–1948.

16. Schumacher, C. and Shima, T., "Single-Task Tours," UAV Cooperative Decision and Control, edited by T. Shima and S. Rasmussen, SIAM, 2009, pp. 15–36.

17. Sylvester, C. A.,Wiens, G. J., and Fitz-Coy, N. G., "Control of collaborative mobile robots subject to nonholonomic constraints," Proceedings of the International Symposium on Collaborative Technologies and Systems, 2004.

18. Santilli, C., Bicchi, A., Casalino, G., and Balestrino, A., "Nonholonomic, bounded curvature path planning in cluttered environments," Proceedings of the Conference on Emerging Technologies and Factory Automation, Vol. 2, 1995, pp. 363–372.

19. Bicchi, A., Casalino, G., and Santilli, C., “Planning Shortest Bounded-Curvature Paths for a Class of Nonholonomic Vehicles among Obstacles,” Proceedings of the IEEE International Conference on Robotics and Automation, 1995, pp. 1349–1354.

20. Jiang, K., Seneviratne, L. D., and Earles, S.W. E., “A shortest path based path planning algorithm for nonholonomic mobile robots,” Journal of Intelligent and Robotic Systems, Vol. 24, 1999, pp. 347–366.

21. Schouwenaars, T., Moor, B. D., Feron, E., and How, J., “Mixed-integer programming for multi-vehicle path planning,” Proceedings of the European Control Conference, 2001, pp. 2603–2608.

22. Kingston, D. and Beard, R., “UAV splay state configuration for moving targets in wind,” Advances in Cooperative Control and Optimization, edited by M. J. Hirsch, P. M. Pardalos, R. Murphey, and D. Grundel, Springer, Inc., 2007, pp. 109–129.

23. Gu, G., Chandler, P. R., Schumacher, C. J., Sparks, A., and Pachter,M., “Optimal cooperative sensing using a team of UAVs,” IEEE Transactions on Aerospace and Electronic Systems, Vol. 42, No. 4, 2006, pp. 1446–1458.

24. Sinclair, A. J., Prazenica, R. J., and Jeffcoat, D. E., “Simultaneous Localization and Planning for Cooperative Air Munitions,” Advances in Cooperative Control and Optimization, edited by M. J. Hirsch, P. M. Pardalos, R. Murphey, and D. Grundel, Springer, Inc., 2007, pp. 81–94.

25. Doucette, E. A., Sinclair, A. J., and Jeffcoat, D. E., “Simultaneous Localization and Planning for Cooperative Air Munitions via Dynamic Programming,” Optimization and Cooperative Control Strategies, edited by M. J. Hirsch, C. W. Commander, P. M. Pardalos, and R. Murphey, Springer, Inc., 2009, pp. 69–79.

26. McLain, T. W., Chandler, P. R., and Pachter,M., “A decomposition strategy for optimal coordination of unmanned air vehicles,” Proceedings of American Control Conference, 2000, pp. 369–373.

27. Rafi, F., Khan, S., Shafiq, K., and Shah,M., “Autonomous target following by unmanned aerial vehicles,” Proceedings of SPIE: Unmanned Systems Technology VIII, edited by G. R. Gerhart, C. M. Shoemaker, and D. W. Gage, Vol. 6230, 2006, pp. 1–8.

28. Sinha, A., Kirubarajan, T., and Bar-Shalom, Y., “Autonomous ground target tracking by multiple cooperative UAVs,” Proceedings of the IEEE Aerospace Conference, 2005, pp. 1–9.

29. Akselrod, D., Sinha, A., and Kirubarajan, T., “Collaborative distributed sensor management to multitarget tracking using hierarchical Markov decision processes,” Proceedings of SPIE: Signal and data processing of small targets, edited by O. E. Drummond and R. D. Teichgraeber, Vol. 6699, 2007, pp. 1–14.

30. Sinha, A., Kirubarajan, T., and Bar-Shalom, Y., “Optimal cooperative placement of UAVs for ground target tracking with doppler radar,” Proceedings of SPIE: Signal processing, sensor fusion, and target recognition XIII, edited by I. Kadar, Vol. 5429, 2004, pp. 95–104.

31. Hirsch, M. J., Ortiz-Pena, H., and Eck, C., “Cooperative Tracking of Multiple Targets by a Team of Autonomous UAVs,” International Journal of Operations Research and Information Systems, Vol. 3, No. 1, 2012, pp. 53–73.

32. Hirsch, M. J., Ortiz-Pena, H., and Sudit, M., “Decentralized Cooperative Urban Tracking of Multiple Ground Targets by a Team of Autonomous UAVs,” Proceedings of the 14th International Conference on Information Fusion, 2011, pp. 1196–1202.

33. Hirsch,M. J. and Ortiz-Pena, H., “Autonomous Network Connectivity and Cooperative Control forMultiple Target Tracking,” Proceedings of the 27th Army Science Conference, 2010.

34. Blackman, S. and Popoli, R., Design and Analysis of Modern Tracking Systems, Artech House, 1999.

35. Gelb, A., Applied Optimal Estimation, MIT Press, 1974.

36. Maybeck, P. S., Stochastic Models, Estimation, and Control, Vol. I, Navtech Press, 1994.

37. Garey, M. R. and Johnson, D. S., Computers and Intractability: A Guide to the Theory of NP-Completeness, W.H. Freeman and Company, 1979.

38. Feo, T. A. and Resende, M. G. C., "A probabilistic heuristic for a computationally difficult set covering problem," Operations Research Letters, Vol. 8, 1989, pp. 67–71.

39. Feo, T. A. and Resende, M. G. C., "Greedy Randomized Adaptive Search Procedures," Journal of Global Optimization, Vol. 6, 1995, pp. 109–133.

40. Festa, P. and Resende, M. G. C., "GRASP: An annotated bibliography," Essays and Surveys in Metaheuristics, edited by C. Ribeiro and P. Hansen, Kluwer Academic Publishers, 2002, pp. 325–367.

41. Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P., "Optimization by simulated annealing," Science, Vol. 220, No. 4598, 1983, pp. 671–680.

42. Matsumoto, M. and Nishimura, T., "Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator," ACM Transactions on Modeling and Computer Simulation, Vol. 8, No. 1, 1998, pp. 3–30.

43. Lerro, D. and Bar-Shalom, Y., "Tracking with Debiased Consistent Converted Measurements Versus EKF," IEEE Transactions on Aerospace and Electronic Systems, Vol. 29, No. 3, 1993, pp. 1015–1022.

44. M.J. Hirsch and D. Schroeder, "Dynamic Decentralized Cooperative Control of Multiple Autonomous Vehicles with Multiple Tasks for Urban Operations," in Proceedings of the AIAA Guidance, Navigation, and Control conference, 2012.

45. T. Roughgarden, Selfish routing and the price of anarchy, MIT Press, 2005.

46. Defense Science Board Task Force Report, "The role of autonomy in DoD systems," 2012.